

AMENDMENTS TO THE CLAIMS

1. (Currently amended) A method of managing communications between service components in a computing environment, the computing environment comprising an interconnection system, an entity external to the interconnection system and communicatively linked with the interconnection system, and a plurality of processing nodes interconnected by the interconnection system, each of the service components having a respective identity ~~and being programmed on at least a respective one of the processing nodes~~, the method comprising:

assigning to each service component a respective trustworthiness measure and a respective criticality measure;

for each service component, using the trustworthiness and criticality measures assigned to the service component so as to determine one or more of the processing nodes onto which the service component should be programmed;

programming each service component onto the one or more processing nodes determined to be the one or more processing nodes onto which the service component should be programmed;

based on the identity of at least one of the service components, establishing access-control logic restricting inter-node communication involving the at least one service component;

programming the external entity with the access-control logic;

at the external entity, receiving from the interconnection system a signal indicating detection of an attempted inter-node communication involving the at least one service component;

in response to receiving the signal, the external entity providing at least a portion of the access-control logic to the interconnection system; and

applying the access-control logic provided to the interconnection system, to block the attempted inter-node communication involving the at least one service component,

wherein the trustworthiness measure for each service component represents an assessment of a potential threat the service component poses to other objects, and

wherein the criticality measure for each service component represents a measure of concern for what the other objects may do to the service component.

2. (Original) The method of claim 1, wherein establishing the access-control logic comprises:

establishing a rule indicating whether to allow a communication involving the at least one service component; and

translating the rule into the access-control logic.

3. (Original) The method of claim 2, wherein:

establishing a rule indicating whether to allow a communication involving the at least one service component comprises establishing a rule indicating whether to allow a communication with a first service component programmed on a first processing node in the computing environment.

4. (Original) The method of claim 3, wherein:

establishing a rule indicating whether to allow a communication involving the at least one service component comprises establishing a rule indicating whether to allow a communication between (i) a first service component programmed on a first processing node in the computing environment and (ii) a second service component programmed on a second processing node in the computing environment.

5. (Previously presented) The method of claim 3, wherein the at least one service component resides at at least one service-access-point in the computing environment, and wherein translating the rule into the access-control logic comprises mapping the rule into packet-filter logic associated with the at least one service-access-point.

6. (Original) The method of claim 5, wherein the at least one service-access-point comprises an IP address of the first processing node.

7. (Original) The method of claim 6, wherein the first processing node is programmed to associate a first transport port with the first service component, and wherein the at least one service-access-point further comprises the first transport port.

8. (Original) The method of claim 1, wherein applying the access-control logic to block an inter-node communication involving the at least one service component comprises:

detecting an attempted inter-node communication involving the at least one service component;

based on the access-control logic, making a determination that the attempted inter-node communication should be blocked; and

in response to the determination, blocking the attempted inter-node communication.

9. (Original) The method of claim 1, wherein at least two processing nodes of the plurality of interconnected processing nodes run different operating systems.

10. (Original) The method of claim 1, wherein at least two processing nodes of the plurality of interconnected processing nodes support different processor instructions sets.

11. (Original) The method of claim 1, wherein the computing environment is a cluster-based computing environment.

12. (Original) The method of claim 1, wherein the computing environment is a public computing platform.

13. (Currently amended) A method of managing communications between service components in a computing environment, the computing environment comprising an interconnection system, an entity external to the interconnection system and communicatively linked with the interconnection system, and a plurality of processing nodes interconnected by the interconnection system, each of the service components having a respective identity ~~and being programmed on at least a respective one of the processing nodes~~, the method comprising:

assigning to each service component a respective trustworthiness measure and a respective criticality measure;

for each service component, using the trustworthiness and criticality measures assigned to the service component so as to determine one or more of the processing nodes onto which the service component should be programmed;

programming each service component onto the one or more processing nodes determined to be the one or more processing nodes onto which the service component should be programmed;

based on the identity of at least one of the service components, establishing at least one access-control rule indicating whether to allow at least one communication involving the at least one service component;

translating the at least one access-control rule into access-control logic;

programming the external entity with the access-control logic;

at the interconnection system, detecting an attempted inter-node communication between service components and responsively sending the external entity a signal indicating detection of the attempted inter-node communication;

at the external entity, receiving the signal indicating detection of the attempted inter-node communication and responsively providing at least a portion of the access-control logic to the interconnection system;

based on the access-control logic provided to the interconnection system, determining that the attempted inter-node communication between service components is not allowed; and
responsively blocking the attempted inter-node communication,

wherein the trustworthiness measure for each service component represents an assessment of a potential threat the service component poses to other objects, and

wherein the criticality measure for each service component represents a measure of concern for what the other objects may do to the service component.

14. (Previously presented) The method of claim 13, wherein the access-control logic comprises packet-filter logic.

15. (Previously presented) The method of claim 13, wherein:
each of the service components is designated by a respective service-access-point (SAP) in the computing environment, and each processing node has a respective SAP as well; and
the access-control logic comprises packet-filter logic associated with at least one SAP in the computing environment.

16. (Previously presented) The method of claim 15, wherein the respective SAP of each service component comprises an IP address of the respective processing node on which the service component is programmed, and wherein the access-control logic comprises packet-filter logic associated with at least one such IP address.

17. (Previously presented) The method of claim 16, wherein at least one of the SAPs of a service component further comprises a port selected from the group consisting of a TCP port and a UDP port, and wherein the packet-filter logic is further associated with at least one such port.

18. (Previously presented) The method of claim 13,
wherein the communications between service components are packet-based; and
wherein the access-control logic comprises packet-filter logic associated with a
combination of at least (i) a packet transport protocol, (ii) a source address in the computing
environment and (iii) a destination address in the computing environment.

19. (Previously presented) The method of claim 13, wherein translating the at
least one access-control rule into access-control logic comprises:
mapping the at least one access-control rule to packet-filter logic associated with at least
one service-access-point in the computing environment.

20. (Original) The method of claim 13,
wherein at least a given one of the processing nodes includes a firewall for restricting
communications with the given processing node; and
translating the at least one access-control rule into access-control logic comprises
provisioning the firewall of the given processing node to allow communications between at least
one service component programmed on the given processing node and at least one service
component programmed on another processing node.

21. (Previously presented) The method of claim 13, wherein the
interconnection system further performs the element of blocking the attempted inter-node
communication.

22-23. (Cancelled)

24. (Previously presented) The method of claim 13, further comprising:
providing at least another portion of the access-control logic to the interconnection system prior to detecting the attempted inter-node communication between service components.

25. (Cancelled)

26. (Previously presented) The method of claim 13, wherein the interconnection system comprises a switch, and wherein providing the at least a portion of the access-control logic to the interconnection system comprises setting up the switch to apply the access-control logic.

27. (Previously presented) The method of claim 26,
wherein the switch comprises (i) a packet-filtering agent and (ii) a provisioning-interface for receiving command-line instructions to set up the packet-filtering agent, the switch being arranged to translate the command-line instructions into packet-filtering logic executable by the packet-filtering agent; and

wherein setting up the switch to apply the access-control logic comprises providing the switch, via the provisioning-interface, with command-line instructions representative of the access-control logic.

28. (Previously presented) The method of claim 21, wherein an entity coupled to the interconnection system performs the element of providing at least a portion of the access-control logic to the interconnection system.

29. (Cancelled)

30. (Previously Presented) The method of claim 28, wherein the external entity coupled to the interconnection system comprises a session manager.

31. (Original) The method of claim 21, wherein the interconnection system comprises a switch.

32. (Original) The method of claim 21, wherein the interconnection system comprises a router.

33. (Original) The method of claim 13, wherein the computing environment is a cluster-based computing environment.

34. (Previously presented) The method of claim 13, wherein the computing environment is a public computing platform.

35. (Original) The method of claim 13, wherein the attempted inter-node communication comprises an attempted inter-node communication between antagonistic service components.

36. (Original) The method of claim 13, wherein the attempted inter-node communication comprises an attempted communication of a packet from a first processing node to a second processing node, and wherein blocking the attempted communication comprises dropping the packet.

37. (Currently amended) A method for managing application logic in a public computing platform, the public computing platform comprising (i) a network of processing nodes interconnected by an interconnection system, and (ii) an entity external to the interconnection system and communicatively linked with the interconnection system, the method comprising:

receiving specifications of at least two computer-program applications, the applications cooperatively comprising a number of application components;

assigning to each application component a respective trustworthiness measure and a respective criticality measure;

for each application component, using the trustworthiness and criticality measures of the application component so as to determine one or more of the processing nodes onto which the application component should be programmed;

programming each application component onto the one or more processing nodes determined to be the one or more processing nodes onto which the application component should be programmed;

~~loading the application components of the at least two applications onto at least two of the processing nodes of the computing platform;~~

at the interconnection system, detecting an attempted inter-node communication between the application components and responsively sending to the external entity a signal indicating detection of the attempted inter-node communication;

at the external entity, receiving the signal indicating detection of the attempted inter-node communication and responsively providing to the interconnection system at least a portion of access-control rules that define allowed communications between the application components; and

applying the access-control rules provided to the interconnection system, to block the attempted inter-node communication between the application components,

wherein the trustworthiness measure for each application component represents an assessment of a potential threat the application component poses to other objects, and

wherein the criticality measure for each application component represents a measure of concern for what the other objects may do to the application component.

38. (Currently amended) A computing environment with communication control comprising:

an interconnection system;

a plurality of co-located processing nodes interconnected via the interconnection system;

a plurality of application components, each of the application components being assigned a respective trustworthiness measure and a respective criticality measure, wherein for each application component, the trustworthiness and criticality measures assigned to the application component are used to determine one or more processing nodes onto which the application component should be loaded, each of the application components being loaded onto the one or more the processing nodes determined to be the one or more processing nodes onto which the application component should be loaded, each of the application components having a respective service-access-point defining a location of the application component in the computing environment; and

an entity, external to the interconnection system, communicatively linked with the interconnection system,

wherein the external entity includes access-control logic indicating allowed inter-node communications between application components,

wherein the interconnection system detects an attempted inter-node communication between application components and responsively sends to the external entity a signal indicating detection of the attempted inter-node communication,

wherein the external entity receives the signal and responsively provides the access-control logic to the interconnection system,

the access-control logic being executable, in response to the attempted inter-node communication, to make a determination of whether the attempted inter-node communication is allowed; and

the access-control logic being executable, in response to a determination that the attempted inter-node communication is not allowed, to block the attempted inter-node communication;~~and,~~

wherein the trustworthiness measure for each application component represents an assessment of a potential threat the application component poses to other objects, and

wherein the criticality measure for each application component represents a measure of concern for what the other objects may do to the application component.

39-41. (Cancelled)

42. (Original) The computing environment of claim 38, wherein the computing environment is a cluster-based computing environment.

43. (Original) The computing environment of claim 38, wherein the computing environment is a public-computing platform.

44. (Original) The method of claim 38, wherein at least two of the co-located processing nodes run different operating systems.

45. (Original) The method of claim 38, wherein at least two of the co-located processing nodes support different processor instructions sets.

46-53. (Cancelled)

54. (New) The method of claim 1, further comprising:

using the access-control logic to block a number of attempted inter-node communications involving the at least one service component, wherein the attempted inter-node communication originates from a given service component;

using the access-control logic to log the number of blocked attempted inter-node communications involving the at least one service component and originating from the given service component, and

after a threshold number of attempted inter-node communications involving the at least one service component and originating from the given service component have been blocked, using the access-control logic to send an alert message to the external entity,

wherein in response to the alert message, the external entity terminates the given service component.

55. (New) The method of claim 54, wherein the threshold number of attempted inter-node communications involving the at least one service component have been blocked is greater than or equal to one.

56. (New) A method of managing communications between service components in a computing environment, the computing environment comprising an interconnection system, an entity external to the interconnection system and communicatively linked with the interconnection system, and a plurality of processing nodes interconnected by the interconnection system, each of the service components having a respective identity, and each of

the service components being programmed on at least a respective one of the processing nodes,
the method comprising:

based on an identity of a first service component and an identity of a second service component, establishing access-control logic to restrict inter-node communications, wherein the inter-node communications involve the first service component and originate from the second service component;

programming the external entity with the access-control logic;

receiving at the external entity, a signal sent from the interconnection system, wherein the signal indicates the interconnection system has detected an attempted inter-node communication, and wherein the attempted inter-node communication originates from the second service component and involves the first service component;

in response to receiving the signal, the external entity providing the access-control logic to the interconnection system;

at the interconnection system, applying the provided access-control logic so as to block the attempted inter-node communication and to log a number of blocked communications, wherein blocking the attempted inter-node communication causes the logged number of blocked communications to reach a threshold number; and

in response to the number of blocked communications reaching the threshold number, applying the provided access-control logic to send an alert message to the external entity,

wherein in response to the alert message, the external entity terminates the second service component.